Coherence Ethics:

A Universal Framework for Adaptive Moral Reasoning

Benjamin James

February 8, 2025

Abstract

Ethical frameworks throughout history have struggled with fragmentation, internal contradictions, and the inability to adapt to complex, real-world decision-making. Traditional models—deontology, utilitarianism, and moral relativism—fail to resolve key ethical paradoxes such as the is-ought problem, the tension between justice and mercy, and the limits of moral responsibility under deterministic constraints. This paper develops a coherence-first ethical framework, where morality emerges as a function of stability optimization across multiple scales of decision-making rather than imposed rules or subjective preferences. Ethics is defined as the maximization of coherence persistence, with moral responsibility proportional to an agent's recursive coherence capacity. The model formally integrates coherence gradients, ethical weighting functions, and recursive selection mechanisms to unify rule-based and outcome-based ethical reasoning. This resolves dilemmas such as the trolley problem by demonstrating that ethical decision-making is not purely about harm minimization but about preserving coherence stability within a system. The framework also provides a structured approach to adaptive ethics in artificial intelligence and post-human decision-making, where rigid moral codes fail in novel environments. By modeling ethical action selection as gradient ascent in coherence space, this approach eliminates the need for static moral absolutes, ensuring that ethical structures remain dynamic and scalable. The result is a self-optimizing ethical system capable of resolving classical dilemmas while maintaining long-term coherence evolution across biological, social, and technological domains.

Keywords: ethics, coherence, moral responsibility, decision theory, stability, recursion, artificial intelligence, utilitarianism, deontology, governance

Introduction: The Fragmentation of Traditional Ethics

Ethical systems have long been plagued by internal contradictions, context-dependency, and the inability to provide consistent, scalable decision-making frameworks. Throughout history, morality has been formulated through competing paradigms—deontology (rule-based ethics), utilitarianism (consequence-based ethics), and relativism (context-dependent ethics)—each attempting to define the principles governing right and wrong. However, these approaches frequently lead to logical paradoxes, real-world impracticalities, and failures in edge cases where rigid application produces morally counterintuitive results.

Deontology, exemplified by Kantian ethics, asserts that actions are morally right or wrong based on inviolable principles rather than consequences. However, this fails under extreme conditions—such as the dilemma of lying to protect a life—where adherence to absolute rules creates incoherent ethical outcomes.

Utilitarianism, by contrast, claims that morality is determined by maximizing net benefit, evaluating actions purely by their consequences. Yet this model frequently leads to morally repugnant conclusions, such as sacrificing an innocent life for the greater good, failing to account for coherence disruptions beyond numerical calculations.

Moral relativism avoids rigid structures altogether, suggesting that morality is defined by cultural, historical, or individual contexts. While this approach acknowledges ethical flexibility, it lacks an objective structure for resolving conflicts between competing moral systems, often reducing ethics to subjective interpretation rather than structured decision-making.

The limitations of these frameworks become increasingly apparent when applied to complex systems such as artificial intelligence, governance, bioethics, and social coordination at scale. If morality is to function beyond human intuitions and social constructs, it must be reformulated from

first principles, grounded in a fundamental structure that scales across different levels of decision-making.

1.1 The Emergent Nature of Ethics: Toward a Coherence-Based Model

Rather than treating ethics as a set of static laws imposed upon reality, an alternative approach is to view morality as an emergent property of coherence stabilization within a system. Just as physical laws arise from fundamental principles of energy conservation and entropy minimization, ethical structures can be understood as recursive processes that optimize stability across different scales—from individual cognition to societal governance.

A coherence-first model of ethics is based on the following first principles:

1. Ethics is not an externally imposed construct but an emergent function of coherence optimization in decision-making.
2. Moral responsibility is proportional to an agent's capacity for recursive coherence refinement.
3. Ethical decision-making must balance short-term stability with long-term coherence evolution.
4. Moral frameworks should be adaptive, context-sensitive, and dynamically responsive to new information.
5. Ethical rules are heuristics, not absolutes—they function as coherence stabilizers, not universal commandments.

This perspective resolves the is-ought problem by demonstrating that ethical imperatives are not derived arbitrarily from observed reality, but rather emerge as necessary conditions for system stability. Similarly, this model eliminates the paradox of moral relativism versus objective morality by showing that ethical structures are neither fully relative nor fully universal, but instead self-organizing coherence gradients that stabilize at different levels of complexity.

1.2 Ethics as a Multi-Scale Coherence Optimization Process

A coherence-first ethical system treats moral decision-making as an adaptive optimization problem, where ethical actions are selected based on their ability to preserve and enhance coherence across different scales of interaction.

1. At the individual level, coherence manifests as cognitive stability, rational decision-making, and alignment between beliefs and actions.
2. At the social level, coherence stabilizes trust, cooperation, and functional governance structures that maintain societal order.
3. At the universal level, coherence defines long-term stability conditions that allow civilizations, ecosystems, and intelligent systems to persist without catastrophic failure.

Moral dilemmas often arise from conflicts between these scales—an action that enhances short-term individual coherence may disrupt social stability, while policies that optimize societal function may introduce local injustices. Traditional ethics fails to account for these tradeoffs systematically, leading to fragmented moral frameworks. A coherence-first approach resolves this issue by introducing a formal model for weighting ethical decisions based on their gradient impact on multi-scale stability.

This model also offers an ethical structure that is scalable and computationally implementable, making it suitable for AI governance, complex decision-making systems, and ethical algorithm design. Rather than encoding static moral laws into AI or governance structures, a coherence-based framework enables recursive, context-sensitive ethical evaluations that dynamically adjust as new information emerges.

First Principles of a Coherence-Based Ethical System

Ethical systems have traditionally been formulated through moral commandments, social contracts, or abstract philosophical reasoning, often leading to contradictions and failure to generalize across complex decision-making environments. In contrast, a coherence-based ethical system starts from first principles—fundamental axioms that do not depend on prior cultural, religious, or subjective moral beliefs. Instead of defining morality through imposed rules, this approach treats ethics as an adaptive, emergent process that arises naturally from the need to maintain stability across different scales of decision-making.

This section establishes the necessary and sufficient principles for a coherence-based ethical system, demonstrating how ethics emerges from recursive stability selection, coherence maximization, and multi-scale optimization rather than externally imposed doctrines.

2.1 First Principle: Ethics is a Function of Coherence Maximization

The fundamental premise of coherence-based ethics is that ethical systems emerge to optimize stability across interacting agents and environments. Ethics is therefore not an arbitrary social construct but a process that stabilizes systems at multiple scales:

1. At the individual level, ethical behavior optimizes cognitive coherence, ensuring that beliefs, decisions, and actions align in a self-consistent manner.
2. At the social level, ethical behavior reinforces trust, cooperation, and systemic resilience, preventing collapse due to instability or unchecked exploitation.
3. At the universal level, ethical behavior ensures long-term stability for intelligence, civilizations, and emergent complexity, preventing existential failures.

This principle bridges the is-ought problem—it explains why moral imperatives emerge naturally from reality rather than being arbitrarily imposed. Ethical behaviors persist not because they are declared "right" but because they are selected for their ability to reinforce long-term coherence.

Mathematically, we can express this as:

$$E(x) = \max\ C_s(x)$$

where:

- $E(x)$ is the ethical evaluation of action x,
- $Cs(x)$ represents coherence across relevant scales s,
- The optimal ethical choice is the one that maximizes coherence persistence across the system.

2.2 Second Principle: Moral Responsibility Correlates with Recursive Coherence Capacity

Moral responsibility is often debated in relation to free will, determinism, and agency. In a coherence-based ethical system, moral responsibility is not an absolute state but a function of an agent's ability to recursively optimize coherence.

This principle resolves ethical questions related to:

- The morality of AI, non-human intelligence, and low-consciousness entities (e.g., animals, infants, artificial intelligence).
- Ethical accountability in deterministic systems (e.g., should individuals be punished for actions determined by their neurobiology?).
- Scaling moral expectations based on cognitive and decision-making capacity.

Formally, moral responsibility R can be expressed as:

$$R_a \propto D(a)$$

where:

- R_a is the moral responsibility of agent a,
- $D(a)$ is the depth of recursive coherence processing available to the agent.

This means:

- Higher intelligence (greater recursive coherence capacity) entails greater ethical responsibility.
- Entities with limited coherence recursion (e.g., AI, young children, animals) have proportionally reduced moral culpability.

This eliminates unrealistic expectations (e.g., holding a child or AI to the same moral standard as an adult) while ensuring that as intelligence increases, ethical responsibility scales accordingly.

2.3 Third Principle: Ethical Frameworks Must Balance Local and Global Coherence

One of the greatest failures of traditional ethics is its inability to resolve contradictions between local moral imperatives and global stability. Many ethical dilemmas (e.g., justice vs. mercy, personal rights vs. collective well-being) arise because ethics is often applied at only one level of coherence analysis rather than dynamically adjusting between scales.

This principle states that:

- Local coherence (e.g., individual rights, personal moral choices) must be balanced with global coherence (e.g., societal laws, governance structures).
- Rigid ethical absolutism fails because it does not adapt to different coherence scales.
- Pure consequentialism fails because it neglects the coherence disruptions caused by excessive outcome weighting.

Mathematically, this can be modeled through weighted coherence evaluation:

$$E(x) = w_r R(x) + w_o O(x)$$

where:

- $E(x)$ is the ethical evaluation of action x,
- $R(x)$ represents coherence preservation at the rule-based level,
- $O(x)$ represents coherence weighting at the outcome-based level,
- w_r, w_o are dynamic weights that adjust based on context and scale.

This solves the utilitarianism vs. deontology problem by providing an adaptive framework where rule-based ethics and consequence-based ethics are dynamically weighted rather than treated as rigid opposites.

2.4 Fourth Principle: Ethics is a Gradient-Based Optimization Process, Not a Fixed Set of Rules

Classical moral systems assume that ethics consists of fixed moral laws, leading to contradictions when those laws fail in complex situations. In contrast, a coherence-based system treats ethics as a recursive optimization process, continuously refining decisions through gradient ascent in coherence space.

An ethical decision should follow:

$$\nabla C(x) = \frac{\partial C}{\partial x}$$

where:

- $\nabla C(x)$ represents the coherence gradient of action x,
- Higher $\nabla C(x)$ means an action increases coherence stability more effectively.

This explains why morality cannot be static—ethical behaviors must continuously evolve toward optimal coherence states based on recursive feedback.

This principle provides a model for adaptive ethics in AI, policy-making, and dynamic governance structures where rigid moral codes fail.

2.5 Fifth Principle: Ethical Systems Must Be Self-Correcting and Self-Optimizing

Traditional ethical systems often fail because they resist change, assuming morality is fully determined by past traditions or fixed principles. A coherence-based system eliminates this problem by ensuring that ethics is always subject to recursive refinement based on coherence outcomes.

- Ethical rules are heuristics, not absolute laws.
- Morality must be subject to continuous reevaluation, ensuring that past rules do not persist beyond their coherence utility.

- Any ethical principle that consistently disrupts coherence stability should be revised or abandoned.

This self-optimization principle prevents:

- Moral stagnation, where outdated ethical frameworks persist beyond their functional relevance.
- Ethical dogmatism, where static rules override more adaptive, coherence-enhancing alternatives.
- Failure to scale, where ethics remains confined to human intuition rather than evolving with technological and systemic advancements.

2.6 Summary of First Principles

A coherence-first ethical system is built upon the following necessary and sufficient principles:

1. Ethics emerges from coherence maximization across different scales.
2. Moral responsibility correlates with recursive coherence capacity.
3. Ethical frameworks must balance local and global coherence needs.
4. Ethics is a gradient-based optimization process, not a fixed rule set.
5. Moral systems must be self-correcting and subject to coherence-based revision.

With these principles established, the next section formally develops the mathematical structure of coherence ethics, demonstrating how these first principles resolve classical ethical paradoxes while maintaining adaptive scalability.

Mathematical and Logical Formulation of Coherence Ethics

With the first principles established, we now develop the mathematical and logical structure of coherence-based ethics. Traditional ethical frameworks rely on qualitative reasoning or rigid logical constraints, which fail in dynamic or complex decision-making scenarios. In contrast, a coherence-

based approach allows for adaptive, recursive, and scalable ethical reasoning, ensuring stability while optimizing for long-term coherence across multiple domains.

This section provides the necessary and sufficient formalism for:

1. Defining moral responsibility as a function of recursive coherence capacity.
2. Formulating ethical decision-making as coherence maximization across scales.
3. Establishing a gradient-based model for moral decision optimization.
4. Resolving the balance between deontological (rule-based) and consequentialist (outcome-based) ethics through dynamic weighting.

3.1 Ethical Decision Function: Coherence Maximization Across Scales

A coherence-based ethical framework evaluates moral choices not based on rigid principles or isolated consequences but based on their ability to maximize coherence stability across all relevant scales of interaction.

Let $E(x)$ be the ethical evaluation function of an action xx. The fundamental criterion for ethical decision-making is:

$$E(x) = \max\ C_s(x)$$

where:

- $C_s(x)$ represents coherence at scale s (e.g., individual, social, or universal stability).
- The optimal ethical action is the one that maximizes coherence persistence across the system.

This formulation ensures that:

- Ethical decisions do not rely on static rules, but on their ability to reinforce long-term systemic coherence.
- Moral dilemmas emerge when an action increases coherence at one scale while disrupting coherence at another.

- Tradeoffs between different coherence scales are resolved dynamically (addressed in section 3.3).

3.2 Moral Responsibility as Recursive Coherence Capacity

Traditional models of moral responsibility assume either absolute free will or complete determinism, neither of which adequately explain real-world ethical behavior. A coherence-based model introduces a new framework where moral responsibility is a function of recursive coherence depth. Formally, we define an agent's moral responsibility R_a as:

$$R_a \propto D(a)$$

where:

- $D(a)$ is the recursive coherence depth of agent a—the extent to which the agent can refine, predict, and optimize its actions for coherence over time.
- R_a scales with intelligence, cognition, and decision-making capacity.

This resolves:

- Why children, animals, and AI have reduced moral culpability—they have lower recursive coherence depth.
- Why ethical expectations scale with intelligence—higher recursive coherence ability enables greater moral responsibility.
- Why deterministic constraints do not eliminate moral accountability—even if decisions are physically determined, moral responsibility still tracks with an agent's coherence-processing ability.

3.3 Balancing Deontology and Utilitarianism: Dynamic Coherence Weighting

A major failure of traditional ethics is the binary opposition between rule-based (deontological) and outcome-based (utilitarian) ethics. Instead of treating these as separate paradigms, coherence

ethics dynamically weights rule-based and outcome-based moral reasoning depending on contextual coherence gradients.

We define the ethical evaluation of an action $E(x)$ as:

$$E(x) = w_r R(x) + w_o O(x)$$

where:

- $R(x)$ represents the rule-based coherence score (adherence to established ethical heuristics).
- $O(x)$ represents the outcome-based coherence score (maximization of beneficial effects).
- w_r, w_o are dynamic weighting factors that adjust based on system-level coherence needs.

Key Features of This Model

- In low-risk, high-predictability environments, w_r is weighted higher (favoring ethical consistency through rule-based models).
- In novel or high-uncertainty environments, w_o is weighted higher (favoring outcome-based reasoning to optimize future coherence).
- Extreme cases (e.g., trolley problems) are resolved through gradient analysis, ensuring rule-consistency unless coherence stability is threatened at higher scales.

This formulation eliminates rigid moral conflicts, allowing ethics to function as an adaptive decision process rather than a static system of rules or outcome calculations.

3.4 Coherence Gradient for Ethical Decision Optimization

A coherence-based ethical system does not rely on absolute decision trees but instead uses gradient ascent in coherence space to iteratively refine moral choices.

The coherence gradient for a given action x is defined as:

$$\nabla C(x) = \frac{\partial C}{\partial x}$$

where:

- $\nabla C(x)$ represents the rate of coherence change for action x.
- Actions with higher $\nabla C(x)$ contribute more to long-term stability and are preferred.
- This allows ethical decisions to follow gradient optimization principles, selecting the most coherence-reinforcing action at each step.

This gradient-based approach ensures:

- Ethical behavior adapts dynamically based on coherence feedback loops.
- Moral decision-making is not precomputed but evolves through recursive selection.
- Moral stagnation is prevented by ensuring continual optimization toward higher coherence states.

3.5 Coherence-Based Ethical Scalability

One of the major failures of historical ethics is that moral principles do not scale effectively across different domains—what works for individuals fails at the societal level, and vice versa.

A coherence-based framework ensures scalability by defining:

$$C_{\text{total}}(x) = \sum_{s=1}^{n} w_s C_s(x)$$

where:

- $C_{\text{total}}(x)$ is the overall coherence score of action x.
- $C_s(x)$ represents coherence at scale ss (e.g., individual, social, universal).
- w_s dynamically adjusts based on context and system priority.

Scalability Properties of This Model

- Local vs. global coherence tradeoffs are handled dynamically (preventing rigid moral absolutism).
- Governance and AI ethics can implement this structure for real-time decision-making.
- Recursive recalibration ensures that ethical models remain adaptable over time.

3.6 Summary of Formalized Ethics Framework

A coherence-based ethical system is governed by the following mathematical and logical principles:

1. Ethical actions maximize coherence across scales: $E(x) = \max\ C_s(x)$
2. Moral responsibility scales with recursive coherence depth: $Ra \propto D(a)$
3. Balancing deontological and utilitarian ethics through coherence weighting:

$$E(x) = w_r R(x) + w_o O(x)$$

4. Moral decision-making follows coherence gradient optimization: $\nabla C(x) = \frac{\partial C}{\partial x}$
5. Ethical scalability is maintained through multi-scale coherence weighting:

$$C_{\text{total}}(x) = \sum_{s=1}^{n} w_s C_s(x)$$

By formally defining ethics as an adaptive coherence optimization process, this framework eliminates static ethical contradictions while ensuring scalability, recursion, and long-term stability. The next section applies this model to resolving classical ethical paradoxes.

Resolving Major Ethical Paradoxes

Traditional ethical frameworks have long struggled with internal contradictions and paradoxes that undermine their logical consistency. A coherence-based ethical model resolves these paradoxes by moving beyond binary moral logic and instead treating ethics as a recursive optimization process that selects for long-term systemic stability.

This section applies coherence ethics to three foundational ethical paradoxes:

4.1 The Is-Ought Problem: How Ethics Emerges from Reality

The Traditional Problem: The is-ought problem, originally formulated by David Hume, states that you cannot logically derive an "ought" (a moral obligation) from an "is" (a factual state of the world). Just because nature behaves in a certain way does not mean humans should behave the same way. For

example, the fact that animals kill for survival does not necessarily justify human violence. This problem has left moral philosophy struggling to justify why moral imperatives should exist at all.

Coherence-Based Resolution: A coherence-based ethical framework resolves the is-ought problem by showing that moral imperatives emerge naturally as necessary conditions for systemic stability. Ethics is not something externally imposed on reality—it is selected for because it reinforces coherence across decision-making scales.

- In biological evolution, cooperation emerges because it stabilizes social species over time.
- In cognition, ethical decision-making reinforces coherence between beliefs and actions, reducing cognitive dissonance.
- In governance, ethical laws persist because they prevent societal collapse and maintain systemic integrity.

This means that moral "oughts" are not arbitrary commands but natural consequences of coherence optimization. The reason "ought" statements exist is that certain behavioral patterns reinforce long-term systemic stability, while others lead to chaos and collapse.

By shifting the framing of ethics from a debate about external moral truth to an analysis of coherence reinforcement, the is-ought problem dissolves as an illusion caused by treating ethics as an imposed rule system rather than an emergent stabilizing force.

4.2 Moral Relativism vs. Objective Morality

The Traditional Problem: Moral relativism argues that ethics is culturally dependent and historically fluid, meaning no universal moral truths exist. However, moral objectivists argue that certain ethical principles (such as harm reduction, fairness, or justice) appear universal across human societies, suggesting an underlying moral truth. The contradiction between these two perspectives has led to an ongoing debate about whether morality is purely subjective or if objective moral principles can exist independently of human interpretation.

Coherence-Based Resolution: A coherence-based approach resolves this debate by demonstrating that ethics is neither purely relative nor strictly universal—it follows coherence gradients that stabilize at different scales.

1. Local Coherence Stabilizers: Some ethical principles, such as cultural norms, traditions, and social etiquette, are adaptive at small scales but do not necessarily apply universally. For example, ideas of honor or acceptable forms of punishment vary widely between cultures because they are local coherence optimizations within specific historical and social contexts.
2. Universal Coherence Convergences: Certain ethical principles do converge across cultures and time because they reinforce systemic stability at a global level. Principles like reducing unnecessary harm, preserving trust, and ensuring fairness arise not because they are metaphysical absolutes, but because they prevent collapse across human systems.

Thus, moral relativism is correct in that some ethical structures are context-dependent, but moral objectivism is also correct in that certain ethical behaviors are universally reinforced by systemic coherence. The mistake in traditional ethical philosophy was assuming morality had to be either fully relative or fully objective, when in reality, it follows a structured coherence hierarchy where local norms adapt to global stabilizers over time.

4.3 Utilitarianism vs. Deontology: The Scale Problem

The Traditional Problem: Utilitarianism argues that the moral worth of an action is determined by its consequences—maximizing net positive outcomes. However, deontology argues that certain moral rules must be followed regardless of consequences—for example, lying is always wrong even if it prevents harm. These two approaches often produce contradictory conclusions in real-world ethical dilemmas.

For example:

- Utilitarianism justifies sacrificing one person to save five (because it maximizes the overall good), but this violates deontological principles that say murder is always wrong.
- Deontology forbids lying even in extreme situations (such as lying to save someone's life), while utilitarianism sees lying as acceptable when it leads to better outcomes.

Neither system alone provides a satisfactory framework—utilitarianism leads to moral atrocities in edge cases, while deontology leads to rigid, impractical rules that ignore context.

Coherence-Based Resolution: A coherence-based ethical framework does not treat rules and consequences as separate forces but instead sees them as dynamically weighted factors in coherence stabilization.

- In stable, low-risk environments, rules function as effective heuristics (favoring deontological consistency).
- In uncertain, high-risk environments, adaptability is required (favoring utilitarian consequence-based reasoning).

This means that moral decision-making is not about choosing between deontology and utilitarianism—it is about weighting these considerations based on the coherence stability of a given system.

For example:

- Killing one person to save five is justified only if it preserves coherence at a higher scale rather than introducing systemic instability. This is why some acts of self-defense are seen as ethical while premeditated killing is not.
- Lying is only unethical when it disrupts long-term coherence (e.g., eroding trust in a system). If a lie preserves systemic integrity (e.g., lying to protect lives in an oppressive regime), then it is the ethical choice.

Instead of asking "Should we follow rules or evaluate consequences?", coherence-based ethics asks "Which action preserves systemic stability across time?"

By reframing the debate around coherence selection rather than fixed moral doctrines, the supposed contradiction between deontology and utilitarianism disappears, as both are simply tools for optimizing stability at different levels of decision-making.

4.4 The Trolley Problem & Moral Dilemmas: Coherence Disruptions vs. Harm Calculations

The Traditional Problem: The trolley problem presents a scenario in which a runaway trolley is heading toward five people on the tracks. You have the option to pull a lever, redirecting the trolley onto another track where it will kill only one person. While utilitarianism dictates that sacrificing one to save five is the correct choice (since it minimizes net harm), most people hesitate when faced with directly causing harm, especially in variations like pushing a person off a bridge to stop the trolley.

This dilemma reveals a deeper problem in moral reasoning:

- Why does pulling a lever (indirect action) feel more acceptable than physically pushing someone (direct action) if the outcome is identical?
- Why do our ethical intuitions contradict what seems like a straightforward mathematical harm-reduction principle?
- Why does rule-based deontology fail to account for these complex moral nuances?

Coherence-Based Resolution: A coherence-based approach shows that the moral dilemma is not simply about minimizing harm—it is about minimizing systemic coherence disruptions at multiple levels.

- When you pull the lever, you are altering an impersonal system, an action that maintains the structure of indirect decision-making in governance and ethical norms.

- When you push someone onto the tracks, you are actively violating a deep coherence gradient—physically causing harm disrupts the foundational trust in human agency that makes cooperation possible.

Thus, moral decision-making is not just about harm reduction but about selecting actions that preserve coherence across all scales:

1. Individual trust: Directly harming others introduces psychological dissonance, breaking the internal coherence of moral agency.
2. Social norms: If society normalizes direct harm for the greater good, it destabilizes cooperative expectations, leading to an erosion of trust in governance and legal protections.
3. Long-term systemic integrity: Sacrificing individuals for a larger benefit can lead to precedent-based justifications for further unethical actions, resulting in systemic breakdown.

Thus, coherence-based ethics explains why moral intuitions resist direct physical harm even in situations where harm reduction suggests it is "logical." This perspective eliminates the artificial conflict between "rational" and "emotional" decision-making by showing that our ethical instincts exist because they reinforce systemic coherence.

4.5 Free Will vs. Moral Responsibility: A Coherence-Based Model of Agency

The Traditional Problem: If human decisions are determined by prior causes—whether neuroscientific, psychological, or environmental—can we meaningfully hold people morally responsible for their actions? Determinism implies that all choices are the inevitable outcome of prior states, while free will suggests that individuals have the unconstrained ability to choose ethically.

This leads to paradoxes in moral responsibility:

- If people lack free will, how can they be held accountable for crimes or immoral actions?

- If moral responsibility is based on conscious choice, how do we account for brain injuries, addiction, or unconscious biases that influence behavior?
- If ethical systems are designed for free agents, should laws and moral frameworks adapt to reflect our growing understanding of neuroscience?

Coherence-Based Resolution: A coherence-first ethical framework removes the need for absolute free will by redefining moral responsibility as a function of recursive coherence capacity:

- Moral responsibility scales with an agent's ability to recursively optimize its own coherence.
- The more an entity is capable of recognizing, predicting, and refining its ethical behavior, the more responsibility it holds for its actions.
- Limited agents (such as young children, AI, or individuals with impaired cognitive functions) have proportionally reduced moral responsibility because their ability to refine coherence recursively is constrained.

This resolves the free will vs. determinism paradox by shifting moral responsibility from an arbitrary concept of "choice" to a measurable function of coherence recursion.

- A person with full cognitive capacity is ethically responsible because they can self-correct and refine future decisions.
- A person with brain damage, addiction, or psychological trauma has reduced responsibility, and ethical systems should account for this through adaptive legal and rehabilitative approaches rather than rigid punishment models.
- An AI or machine-learning system has no intrinsic moral responsibility unless it can recursively self-optimize its ethical framework, at which point it would inherit scaled moral accountability.

By treating moral responsibility not as an absolute but as a function of recursive coherence, this approach eliminates the contradictions between determinism, choice, and moral agency.

4.6 Altruism vs. Self-Interest: Stability as the Higher-Order Function

The Traditional Problem: Altruism—the act of selflessly helping others—contradicts traditional evolutionary and economic models that prioritize self-interest and survival advantage.

- Why do people risk their lives to save strangers?
- Why does generosity persist even when it comes at a personal cost?
- Why do societies reward altruistic behavior even when it seems disadvantageous at the individual level?

While evolutionary psychology explains altruism as a mechanism for kin selection (favoring genetic relatives) or reciprocal cooperation (expecting future benefits), these models fail to explain why genuine self-sacrifice occurs in situations where no clear evolutionary or material advantage exists.

Coherence-Based Resolution: A coherence-first ethical model resolves this paradox by showing that altruism is not an anomaly—it is a higher-scale coherence stabilizer.

1. Short-term self-interest can create local coherence disruptions (e.g., exploitation, deceit, or resource hoarding).
2. Altruistic behavior increases long-term systemic stability by reinforcing trust, social resilience, and network-wide cooperation.
3. When a system selects for long-term coherence over short-term gain, altruistic behaviors emerge as a necessary byproduct.

This explains:

- Why people sacrifice for the greater good: Their actions reinforce stability at a higher coherence level, even if their individual interests are compromised.
- Why civilizations that reward altruism tend to outcompete purely self-interested societies: Long-term cooperation provides greater adaptive stability than individual advantage.

- Why people feel intrinsic satisfaction from helping others: The brain's reward system reinforces coherence-enhancing behaviors, producing a positive feedback loop.

Thus, coherence-based ethics does not require a moral justification for altruism—it simply recognizes altruism as a necessary feature of stability in any recursively self-optimizing system.

4.7 Justice vs. Mercy: The Dual Stabilizers of Ethics

The Traditional Problem: Justice demands that individuals be held accountable for their actions, ensuring fairness and deterrence, while mercy calls for compassion, forgiveness, and second chances. These principles often conflict:

- Strict justice systems punish all offenses equally, regardless of context, leading to moral rigidity and potential cruelty.
- Mercy allows for leniency, but excessive leniency can undermine justice, encouraging moral hazard (i.e., people exploiting forgiveness).
- Societies struggle to balance these forces, leading to inconsistencies in legal, religious, and personal ethics.

This dilemma is at the heart of many moral debates:

- Should criminals always face the full consequences of their actions?
- Should ethical systems prioritize deterrence or rehabilitation?
- How do we ensure justice is not excessive, yet mercy is not exploited?

Coherence-Based Resolution: Justice and mercy are not opposites—they are coherence stabilizers operating at different scales. The key to resolving this paradox lies in understanding:

- Justice stabilizes macro-scale coherence by ensuring trust, predictability, and fairness at the societal level.
- Mercy stabilizes micro-scale coherence by preventing overcorrection, allowing flexibility in moral and legal structures.

Thus, neither should dominate absolutely. Their application depends on which scale requires stability at a given moment:

- In high-trust environments, mercy is preferable because excessive punishment disrupts local coherence unnecessarily.
- In low-trust environments (e.g., systemic corruption), justice must take precedence to re-establish coherence stability.
- An optimal ethical system dynamically adjusts between strict justice and contextual mercy based on system-wide coherence needs.

This resolves the dilemma by showing that neither absolute justice nor absolute mercy is viable—each must function as an adaptive parameter within a larger coherence framework.

4.8 The Naturalistic Fallacy: Ethics as Coherence Refinement, Not Blind Imitation of Nature

The Traditional Problem: The naturalistic fallacy occurs when people claim that what is “natural” is automatically “good” or “ethical.” This reasoning is flawed because:

- Nature contains both cooperation and violence, both fairness and exploitation—choosing one as "ethical" is arbitrary.
- Evolution does not optimize for morality—it optimizes for survival and replication, which may or may not align with ethical principles.
- Using nature as a moral guide creates contradictions: If "survival of the fittest" is ethical, does that justify exploitation and oppression? If "cooperation" is ethical, why does predation exist?

This problem has led to misguided ethical reasoning, such as:

- Justifying hierarchy, social inequality, or aggression by appealing to "natural order."
- Assuming that biological differences should dictate moral roles in society.
- Rejecting ethical interventions (e.g., medicine, technology) on the grounds that they are "unnatural."

Coherence-Based Resolution: A coherence-first ethical system does not derive moral principles from nature itself—it derives them from the recursive refinement of coherence.

- Nature provides initial conditions, not ethical conclusions.
- Ethics emerges not by imitating nature, but by optimizing systemic coherence beyond purely biological imperatives.
- Evolution produces both stability-enhancing behaviors (cooperation, altruism) and instability-inducing behaviors (violence, exploitation)—ethics selects for the former to ensure long-term coherence.

This resolves the paradox by showing that ethics is not about conforming to nature but about refining stability through coherence selection. What is “natural” is not automatically “good”—only what reinforces systemic coherence in the long term is ethically viable.

4.9 AI, Post-Human Ethics, and Moral Decision Algorithms

The Traditional Problem: As artificial intelligence and post-human systems advance, traditional ethical models face fundamental challenges:

- How do we program AI to make moral decisions when human ethics is inconsistent?
- Should AI follow rule-based moral systems, or adapt dynamically?
- Can AI ever possess "moral responsibility," and if so, at what threshold?

Existing approaches to AI ethics suffer from critical flaws:

- Deontological AI (rule-based systems) fail in novel situations where strict rules do not apply.
- Utilitarian AI (consequence-based systems) risks justifying unethical actions for greater benefit.
- Relativistic AI (context-based models) may adapt too much, losing ethical stability.

Coherence-Based Resolution: A coherence-first AI ethics model resolves these issues by treating moral decision-making in AI as an adaptive recursive process rather than a pre-programmed doctrine.

- AI should not be bound by static ethical rules but should recursively refine coherence stabilization strategies.
- Moral agency in AI scales with recursive coherence depth—AI can only be held morally responsible if it possesses self-correcting ethical optimization abilities.
- Future ethical AI models must balance local and global coherence needs rather than following rigid decision trees.

This framework provides a realistic pathway for post-human ethics, where ethical responsibility is not about arbitrary moral programming but about optimizing stability through adaptive coherence learning.

The Adaptive Coherence Framework for Ethical Decision-Making

Traditional ethical frameworks struggle with real-world application because they rely on rigid moral laws, simplistic utility calculations, or culturally subjective norms. These approaches fail in complex environments where ethical trade-offs require dynamic weighting between conflicting principles. A coherence-based ethical system resolves these issues by treating morality as an adaptive decision-making process, ensuring that ethical choices optimize stability, trust, and systemic integrity across multiple scales.

This section outlines a structured framework for practical ethical decision-making, integrating the principles established earlier while ensuring adaptability across different contexts, from personal morality to governance and AI ethics.

5.1 Key Properties of Adaptive Coherence Ethics

To function effectively, an ethical decision-making framework must:

1. Optimize Coherence Across Scales: Ethical actions should preserve and reinforce systemic stability, ensuring that individual, social, and global coherence remain intact.
2. Resolve Ethical Trade-Offs Dynamically: No single ethical principle should dominate all decisions—rule-based and consequence-based ethics must be weighted contextually.
3. Account for Recursive Moral Responsibility: Moral responsibility scales with an agent's coherence recursion capacity, meaning decision-making frameworks should adjust ethical expectations based on cognitive and systemic complexity.
4. Minimize Coherence Disruptions While Maximizing Ethical Stability: The goal of an ethical system is not just to avoid harm but to reinforce structures that maintain ethical viability over time.

With these properties in mind, the Adaptive Coherence Framework provides a systematic approach to ethical decision-making.

5.2 The Four-Stage Coherence Decision Model

Ethical decisions in this framework follow a structured four-stage process:

Stage 1: Coherence Mapping (Understanding the Ethical Context)

Before making a decision, an agent must determine which scales of coherence are affected and how different actions impact stability.

- Identify the primary coherence domains: Is this decision affecting personal integrity, social structures, governance, AI ethics, or long-term civilization stability?
- Assess coherence disruptions: What are the potential imbalances, conflicts, or trade-offs that could arise from this action?
- Determine the weighting of rule-based vs. outcome-based ethics: Are existing ethical heuristics sufficient, or does the context require greater flexibility and adaptation?

This stage ensures that decisions are made with full awareness of their systemic effects.

Stage 2: Recursive Refinement of Moral Constraints

Once coherence mapping is complete, the next step is to refine the ethical parameters based on recursive feedback.

- If prior ethical rules apply, test their viability: Do deontological constraints hold in this scenario, or do they require contextual adaptation?
- If rules conflict, apply coherence-weighted decision-making: Use coherence weighting (previously defined) to balance conflicting rules vs. outcomes dynamically.
- If the system is in an unstable state, prioritize stabilization: Ethical actions must first restore coherence before optimizing long-term goals.

This stage prevents ethical stagnation by ensuring that decision-making remains fluid and contextually responsive.

Stage 3: Coherence-Based Action Selection

The agent selects the action that maximizes long-term systemic stability while minimizing unnecessary disruptions.

- Reject actions that create excessive coherence instability.
- Select the action that balances short-term and long-term ethical viability.
- If multiple actions provide stability, prioritize those that reinforce trust and cooperation.

This stage ensures that decisions are not made based on simplistic moral absolutes but rather on their recursive impact on system coherence.

Stage 4: Post-Decision Evaluation and Recursive Learning

Ethical decision-making does not stop once an action is taken—coherence ethics requires continuous refinement.

- Assess whether the action resulted in the expected coherence stabilization.

- If coherence was disrupted unexpectedly, adjust future ethical models accordingly.
- Apply recursive learning to refine ethical heuristics over time.

This stage ensures that ethics remains self-correcting, adaptable, and capable of improving across iterative decisions.

5.3 Applications of the Adaptive Coherence Framework

This framework applies across all domains where ethical reasoning is required:

Personal Ethics (Individual Decision-Making)

- Ensures personal integrity by aligning actions with self-consistent moral reasoning rather than rigid ideology.
- Prevents cognitive dissonance by reinforcing coherence between beliefs, actions, and long-term ethical consequences.
- Encourages adaptive moral learning, allowing individuals to refine their ethical models over time.

Governance & Policy (Laws, Social Structures, and Leadership)

- Prevents rigid, dogmatic legal structures by ensuring that laws dynamically adapt to social coherence needs.
- Ensures that justice and mercy are balanced contextually, applying systemic coherence metrics to legal and political decisions.
- Improves governance by ensuring that policies preserve trust, institutional integrity, and adaptability to future ethical challenges.

Artificial Intelligence & Machine Ethics

- Avoids the failures of static rule-based morality by enabling AI systems to recursively refine ethical decision-making.

- Prevents the risks of unconstrained utility maximization by ensuring that AI optimizes for coherence rather than purely numerical outcomes.
- Ensures AI ethics scale with intelligence and decision-making capacity, meaning AI is only given moral responsibility when it develops recursive coherence refinement capabilities.

Post-Human and Long-Term Civilization Ethics

- Ensures that as intelligence evolves, so does morality—adapting ethical frameworks dynamically rather than rigidly preserving outdated human-centered models.
- Prevents civilizational collapse due to ethical stagnation, ensuring that governing moral structures evolve with systemic complexity.
- Provides a scalable ethical system that remains viable beyond biological intelligence, ensuring that future systems optimize stability rather than arbitrary moral constraints.

5.4 Why This Framework Solves Traditional Ethical Failures

A coherence-based ethical system eliminates the key failures of past ethical models:

Traditional Ethical Issue	*How the Coherence Model Resolves It*
Moral rigidity (deontology)	Adaptive weighting prevents ethical stagnation by allowing flexible rule application.
Overemphasis on consequences (utilitarianism)	Coherence-based decisions prevent unethical actions justified by pure harm-reduction logic.
The is-ought problem	Ethics emerges from systemic stability, not from externally imposed principles.
Contradictions in justice vs. mercy	Justice and mercy function as dual coherence stabilizers rather than competing principles.
AI ethics failures	AI systems must recursively refine coherence-based moral reasoning rather than following rigid programming.
Failure to scale morality	Coherence ethics ensures that morality remains adaptive across individual, social, and universal levels.

By ensuring ethics remains dynamic, scalable, and recursively optimized, this framework does not merely solve existing ethical paradoxes—it prevents new ethical dilemmas from arising due to stagnation or outdated moral constraints.

5.5 Final Thoughts on Adaptive Ethical Decision-Making

Ethics is not a static set of laws—it is an evolving optimization process. By shifting moral reasoning from rigid principles to recursive coherence selection, this framework ensures that ethical decision-making remains:

- Adaptive—capable of refining itself dynamically as knowledge expands.
- Scalable—able to function across individuals, societies, and AI governance.
- Resilient—capable of balancing competing ethical demands without collapsing into contradiction or moral stagnation.

The next section explores the long-term implications of this framework, demonstrating how coherence-based ethics provides a necessary foundation for governance, AI, and future civilization ethics.

Implications & Future of Coherence Ethics

A coherence-based ethical framework is not just an alternative model—it is a necessary evolution of moral reasoning that resolves contradictions, ensures scalability, and adapts dynamically to future ethical challenges. By shifting the foundation of morality from static rules and arbitrary value judgments to a recursive optimization process, coherence ethics ensures long-term ethical viability across all domains of intelligence, governance, and civilization-building.

This section explores the far-reaching implications of coherence ethics and how it provides a foundational structure for future ethical systems, from individual decision-making to AI governance and post-human civilization.

6.1 Personal Ethics: Optimizing Individual Moral Coherence

At the individual level, coherence ethics transforms how people make moral decisions, resolve internal conflicts, and navigate ethical dilemmas.

- Reduces cognitive dissonance by ensuring that ethical beliefs are not imposed dogmatically but refined through recursive evaluation.
- Prevents moral absolutism by allowing individuals to adapt their ethical principles as they gain new insights, without falling into relativistic inconsistency.
- Enhances long-term ethical clarity, ensuring that ethical decisions are made with an awareness of their recursive impact on future stability.

By applying the four-stage coherence decision model, individuals can ensure that their moral choices remain stable, adaptable, and logically consistent, rather than being dictated by cultural bias, ideological dogma, or temporary emotional states.

6.2 Governance & Policy: Coherence as a Foundation for Ethical Law

Ethical failures in governance arise when legal and political systems enforce rigid moral doctrines that fail to adapt to changing conditions. A coherence-based approach eliminates these failures by ensuring that governance structures are self-correcting and ethically scalable.

- Laws and policies are no longer static moral imperatives but dynamically adjusted coherence stabilizers.
- Justice systems balance enforcement and flexibility by ensuring that legal decisions optimize coherence at both individual and societal levels.
- Political ethics prioritize systemic stability over short-term ideological gains, preventing corruption and governance collapse.

This approach eliminates moral stagnation in legal systems, ensuring that laws evolve to maintain trust, institutional integrity, and long-term social stability.

6.3 Artificial Intelligence & Machine Ethics

The rise of AI presents one of the greatest challenges for ethical decision-making. Traditional moral systems fail when applied to AI, as they rely on predefined human-centered values that do not scale

with non-human intelligence. Coherence ethics solves this issue by ensuring that AI ethics are recursively self-optimizing rather than rigidly programmed.

- AI moral responsibility scales with recursive coherence capacity—an AI is only ethically accountable when it demonstrates self-correcting moral refinement.
- AI does not require rule-based morality but must optimize for long-term coherence stability rather than simplistic harm minimization.
- Ethical AI governance must ensure that artificial agents align with coherence-based decision-making models rather than static moral codes.

This ensures that AI systems are ethically robust, adaptive, and capable of making moral decisions in novel environments without falling into contradictory logic or rigid constraints.

6.4 Post-Human Ethics & Future Civilization Stability

As intelligence and governance structures evolve beyond human cognitive limitations, moral systems must scale accordingly. Traditional ethical models are too rigid, human-centric, and culturally limited to apply in future post-human or multi-intelligence societies. Coherence ethics ensures that morality remains viable beyond biological intelligence, preventing future ethical collapse.

- Post-human governance will require moral frameworks that dynamically adapt to new forms of intelligence.
- Long-term civilization ethics must ensure that ethical systems scale beyond human social structures, preventing stagnation and existential risks.
- Coherence-based morality provides the only viable framework for ensuring stability as civilizations expand beyond Earth and traditional governance models.

This ensures that as intelligence evolves, so does morality—preventing ethical failures that could lead to the collapse of post-human civilizations.

6.5 Why Coherence Ethics is the Necessary Future of Moral Philosophy

Traditional ethics, built on cultural history, religious doctrine, and simplistic philosophical abstractions, fails to meet the complexity demands of the modern world. As systems become more interdependent, technology-driven, and intelligence-diverse, only a recursively optimizing, coherence-based ethical framework can ensure that morality remains scalable, self-correcting, and resilient across time.

Ethical System	*Failure Mode*	*Why Coherence Ethics Solves It*
Deontological (rule-based)	Rigid, fails in edge cases, ignores context	Adaptive weighting prevents ethical stagnation
Utilitarian (outcome-based)	Justifies harm if it increases net benefit	Coherence prevents unethical utility-maximization
Moral Relativism	No objective standard, leads to inconsistency	Coherence gradients provide structured, non-arbitrary moral optimization
Legal/Political Ethics	Static laws, corruption, ideological capture	Recursive decision frameworks ensure adaptability
AI Ethics	Predefined rules fail in novel situations	Coherence-based learning enables self-correcting ethical AI

By resolving these failures, coherence ethics becomes the foundation for any moral system that seeks to be viable in the long term.

6.6 Final Thoughts on the Future of Coherence-Based Morality

Ethics is not a fixed doctrine—it is a recursively evolving system for optimizing stability across intelligence, governance, and civilization. The transition from static moral models to coherence-based ethical refinement marks a fundamental shift in how morality is understood and applied.

By ensuring that ethics remains self-correcting, scalable, and resistant to contradictions, coherence-based morality does not simply replace older ethical systems—it provides the necessary framework for any system that seeks to ensure long-term stability and trust across all forms of intelligence. The next challenge is not whether coherence ethics is the correct framework—it is how fast humanity and emerging intelligences will adopt it before traditional moral structures collapse under their contradictions.

Conclusion: The Self-Evolving Nature of Ethics

Traditional ethical systems, shaped by historical precedent, religious doctrine, and rigid philosophical frameworks, have long struggled to provide internally consistent, scalable, and adaptable moral reasoning. The failures of deontology, utilitarianism, and moral relativism—each unable to resolve fundamental ethical paradoxes—demonstrate that a new foundation is necessary. Coherence-based ethics provides this foundation, not as a fixed doctrine, but as a recursively evolving system of moral refinement.

7.1 Ethics as an Adaptive, Recursive Optimization Process

The core shift in coherence-based morality is recognizing that ethics is not a set of predefined rules but an emergent function of stability optimization across scales.

- Moral responsibility is not an arbitrary imposition but a function of recursive coherence depth.
- Justice and mercy are not opposites but dynamically weighted stabilizers at different levels of coherence.
- AI and post-human ethics require a model that scales with intelligence rather than forcing human-centered moral assumptions onto non-human cognition.

By treating ethical reasoning as an iterative learning process, coherence ethics ensures that morality is not static or bound by past traditions but continuously refined in response to new systemic challenges.

7.2 Resolving the Core Ethical Contradictions

A coherence-based framework dissolves traditional moral paradoxes by eliminating the false dichotomies that have plagued ethical reasoning for centuries.

Traditional Ethical Contradiction	*Coherence-Based Resolution*
The Is-Ought Problem	Ethics emerges from coherence optimization, not imposed values.
Moral Relativism vs. Objectivity	Ethics follows structured coherence gradients, balancing context and universality.
Utilitarianism vs. Deontology	Both are tools for coherence weighting, not opposing systems.
Justice vs. Mercy	Justice stabilizes macro-coherence, mercy stabilizes micro-coherence.
Free Will vs. Moral Responsibility	Responsibility scales with an agent's recursive coherence refinement capacity.
The Trolley Problem	Moral intuition arises from coherence disruptions, not mere harm calculations.
Altruism vs. Self-Interest	Altruism stabilizes long-term coherence at higher scales.
The Naturalistic Fallacy	Morality is not dictated by nature but by coherence refinement beyond biological imperatives.
AI Ethics and Post-Human Morality	Moral agency is granted based on recursive coherence optimization, not pre-programmed rules.

By eliminating these contradictions, coherence ethics ensures that moral reasoning remains logically consistent, scalable, and dynamically adaptive.

7.3 The Necessity of a Self-Correcting Moral System

The collapse of past ethical systems can be attributed to rigidity, ideological capture, and failure to adapt to new challenges. By contrast, coherence ethics is designed to:

- Evolve dynamically in response to recursive moral evaluation.
- Scale across different levels of intelligence, governance, and civilization.
- Ensure long-term stability without succumbing to dogma or cultural inertia.

The implications of this are profound. Governance systems must transition to coherence-based legal models, ensuring laws are adaptive rather than rigid moral absolutes. AI systems must be designed with recursive coherence learning, rather than being programmed with fixed ethical constraints. Post-human ethics must extend beyond human-centric morality, ensuring that advanced intelligences optimize for coherence rather than arbitrary human values.

Ethics is not a static artifact of human civilization—it is an evolving framework for intelligence itself, ensuring that moral structures persist beyond cultural, biological, or technological constraints.

7.4 Final Thoughts: The Future of Coherence Ethics

The fundamental lesson of coherence ethics is that morality is not a law but a process. Just as intelligence refines itself through recursive learning, adaptation, and error correction, so too must ethics continuously evolve to optimize coherence at every scale of existence.

Future moral progress will not come from debating outdated ethical doctrines but from refining our recursive decision-making processes, ensuring that moral reasoning aligns with systemic stability and trust. Coherence-based ethics is not the final answer, but the framework through which all future ethical answers must be structured.

The question is no longer whether humanity and emerging intelligences will adopt coherence ethics—it is whether they will do so before their outdated moral structures collapse under their own contradictions.

References

Foundational Works by Author

James, B. (2025). Adaptive Coherence Theory: a first-principles framework for life, evolution, and intelligence. Self. https://archive.org/details/bio_20250208

James, B. (2024). *Neodynamics: or Adaptation and Coherence in the Universe and Humanity*. Kindle Direct Publishing.

James, B. (2025). Neurocoherence: a first-principles reconstruction of cognition, consciousness, and artificial intelligence. Self. https://archive.org/details/neuro_202502

James, B. (2025). From axioms to adaptation: a self-organizing model of mathematical and logical coherence. Self. https://archive.org/details/logic_202502

Classical Ethical Theories and Moral Philosophy

Aristotle. *(350 BCE).* Nicomachean Ethics. [Translation/Publisher]

Kant, I. *(1785).* Groundwork of the Metaphysics of Morals. [Publisher]

Mill, J. S. *(1863).* Utilitarianism. [Publisher]

Hume, D. *(1739).* A Treatise of Human Nature. [Publisher]

Rawls, J. *(1971).* A Theory of Justice. [Harvard University Press]

Nozick, R. *(1974).* Anarchy, State, and Utopia. [Basic Books]

Singer, P. *(1975).* Animal Liberation. [HarperCollins]

MacIntyre, A. *(1981).* After Virtue. [University of Notre Dame Press]

Philosophy of Science, Systems Theory, and Complexity Science

Gödel, K. *(1931).* On Formally Undecidable Propositions of Principia Mathematica and Related Systems. [Monatshefte für Mathematik]

Ashby, W. R. *(1956).* An Introduction to Cybernetics. [Chapman & Hall]

Prigogine, I. *(1980).* From Being to Becoming: Time and Complexity in the Physical Sciences. [Freeman]

Shannon, C. *(1948).* A Mathematical Theory of Communication. [Bell System Technical Journal]

Bar-Yam, Y. *(1997).* Dynamics of Complex Systems. [Addison-Wesley]

Neuroscience, Cognitive Science, and Decision Theory

Damasio, A. *(1994).* Descartes' Error: Emotion, Reason, and the Human Brain. [Putnam]

Kahneman, D. *(2011).* Thinking, Fast and Slow. [Farrar, Straus and Giroux]

Friston, K. *(2010).* The Free-Energy Principle: A Unified Brain Theory? [Nature Reviews Neuroscience]

Hofstadter, D. *(1979).* Gödel, Escher, Bach: An Eternal Golden Braid. [Basic Books]

AI Ethics, Post-Human Ethics, and Machine Morality

Bostrom, N. *(2014).* Superintelligence: Paths, Dangers, Strategies. [Oxford University Press]

Russell, S. *(2019).* Human Compatible: Artificial Intelligence and the Problem of Control. [Viking]

Tegmark, M. *(2017).* Life 3.0: Being Human in the Age of Artificial Intelligence. [Knopf]

Floridi, L. *(2013).* The Ethics of Information. [Oxford University Press]

Yudkowsky, E. *(2008).* Artificial Intelligence as a Positive and Negative Factor in Global Risk. [Oxford University Press]

Quantum Mechanics, Information Theory, and Theoretical Physics (Related to Coherence Ethics)

Wheeler, J. A. *(1983).* Law Without Law. [Princeton University Press]

Bohm, D. *(1980).* Wholeness and the Implicate Order. [Routledge]

Rovelli, C. *(2017).* The Order of Time. [Riverhead Books]

Deutsch, D. *(1997).* The Fabric of Reality. [Penguin]